

REMARKS

Claims 1-27 were previously pending in this application. By this amendment, Applicant is canceling claims 9, 16, and 24 without prejudice or disclaimer. Claims 1, 4, 10-15, 19, 21, 23, have been amended. As a result, claims 1-8, 10-15, 17-23, and 25-27 are pending for examination with claims 1, 15, 21, and 23 being independent claims. No new matter has been added.

Rejections Under 35 U.S.C. §103

The Office Action rejected claims 1-12 and 15-27 under 35 U.S.C. §103(a) as being unpatentable over Anita Jindal, et al., U.S. Patent No. 6,092,178, (hereinafter “Jindal”) in further view of Sunil K. Srivastava, U.S. Patent Application Publication No. 2005/0149531, (hereinafter “Srivastava”). Applicant has amended claims 1, 4, 10-15, 19, 21, and 23 to further describe and clarify the invention, and in response, Applicant respectfully traverses the rejection and submits the following remarks.

Independent Claim 1

Jindal is directed to a system for responding to a resource request that employs a trigger in association with a network naming service, such as DNS (Domain Name Service), that handles client requests for an application. (Abstract). The trigger comprises a set of executable instructions referenced by a resource record associated with an identifier of the application. (Abstract). In response to a client request concerning the application, the resource record is retrieved and the instructions executed. (Abstract). In one embodiment, information (response time, operational status, number of clients connected to the instance, throughput, status or performance of the servers themselves) is stored on the DNS server or a system that is coupled to the DNS server. (Col. 6, lines 60-62; lines 45-54). Identities of one or more of the preferred servers may also be stored. Thus a trigger, when executed in response to a client request, may retrieve the identity of a pre-selected preferred server or select a preferred server based on the collected information and/or criteria provided in the client request. (Col. 6, lines 62-67). In one embodiment, the information is collected by a load-balancing framework, consisting of multiple executable objects, modules or other collection of executable instructions. (Col 3, lines 61-63). Each instance of an application is thus associated with its own status object(s) and monitor

objects that collect and store information. (please see Col 3, lines 66-67 and Col. 4 lines 1-12).

In summary, Jindal discloses a system for responding to a resource request that performs load-balancing by employing triggers, the triggers either select a predefined preferred server or the triggers perform analysis based on stored data to select a preferred server hosting an application instance. Thus, in response to a request from a client for a service, Jindal performs a load-balancing analysis and then connects the client to a “preferred server.”

In contrast, claim 1 recites a method of providing a service to a client from one of a plurality of servers in a server farm, each of the servers arranged to provide the service to the client, each of the servers being associated with a service address common to all of the servers, and the servers communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm. The method comprises the steps of receiving, at a DNS system, a request for the service from the client, the request specifying the common service address, in response to the request, applying a load balancing method to select a first one of the plurality of servers in the server farm and connecting the client to the first one of the plurality of servers, receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected, and selecting, at the client, a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information.

Jindal fails to teach or suggest servers in a server farm “communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended. Jindal does indicate that if the DNS server does not have access to the type of information needed to chose a preferred server, the DNS server may interactively request the information from one or more servers. (See Col 7, lines 27- 43). However, one skilled in the art would not seek to modify the teachings of Jindal to include storing “identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended. Since client requests are received at the DNS server, storing the information at the DNS server is the obvious choice and storing the information at another server would be considered a disadvantage that would add delay. As Jindal discusses in the background section “requiring the DNS server to query or access another server in order to resolve the request is inefficient and delays satisfaction of the

request.” (Col. 2, lines 11-13). Therefore, Jindal teaches away from the idea of storing information anywhere other than at the DNS server if at all possible, and does not teach or suggest that the information could be stored at a sever in the server farm. Therefore, Jindal could not reasonably teach or suggest storing a copy of the information at *each* of the servers in the servers farm. Jindal teaches using only one data file for storing information collected on the available servers. (Please see e.g. col. 8, lines 15-29). Therefore, Jindal does not teach or suggest servers in a server farm “communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended.

In addition, Jindal does not teach or suggest “receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected” as recited in claim 1, as amended. Jindal describes receiving a client request by a DNS server which retrieves a resource record corresponding to the virtual server name. (Abstract). Within the record is the name of a trigger. (Abstract). The trigger is executed to select, or retrieve an identity of, a sever to which the client request it to be directed. (Abstract). Once Jindal directs the request to the server and the client is connected to the sever that will provide the application, the process described by Jindal ends as a load balanced connection has been established to the application the client has requested. (Please see Abstract). Therefore, Jindal does not teach or suggest “receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected,” as recited in claim 1, as amended. Nor does Jindal teach or suggest “selecting, at the client, a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information,” as recited in claim 1, as amended.

Srivastava does not supply the missing limitations of claim 1. Srivastava is directed to an apparatus and a method for the rapid routing of data to a load-balanced server by employing label values to define a forwarding route, enabling nodes to fast-switch packets based on label mappings. (Abstract). The first server response packet is switched hop-by-hop, and the label is stored at each node traversed by the response packets. (Abstract). For other request and response packets, nodes fast-switch the packets based on the label mappings, thus packet flows are rapidly routed from the client to the same server without time-consuming hop-by-hop routing

or repeated load balancing decisions. (Abstract). In particular, Srivastava does not teach or suggest servers in a server farm “communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended. Thus the combination, assuming without admitting it was proper and feasible, does not render obvious that which is claimed.

In addition, Srivastava does not teach or suggest “receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected,” as recited in claim 1, as amended. Srivastava teaches that after selections are made at load balancing nodes, label records provide a mapping to fast switch subsequent packets. (please see Abstract and para. 0030). Srivastava provides a method for recording load-balancing decisions. Thus it does not teach or suggest “receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected,” as recited in claim 1, as amended. Nor does it teach or suggest “selecting, at the client, a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information,” as recited in claim 1, as amended.

Although Srivastava recites that a client can look up using DNS, available servers for a particular protocol, and can return server preferences that may assist the client in selecting a protocol (para. 5, lines 8-11), it is not taught or suggested that the client may do so at one of the plurality of servers in the server farm where each of the servers are arranged to provide a service to the client. Therefore, the combination of Jindal and the background teachings of Srivastava (assuming, without admitting, that the combination was proper and feasible) would not render obvious that which is recited in claim 1, as amended.

Moreover, Srivastava appears to teach away from the combination proposed by the Examiner. As Srivastava describes in the background, “although this approach is workable, when a plurality of servers is organized in a server farm that is distributed over numerous logically or geographically separate sites, the past approach becomes inefficient.” (para. 0004). Srivastava criticizes conventional approaches of load balancing in the context of high-demand content networks and criticizes past approaches for failing to provide “stickiness.” (See para. 0012 and 0024). Thus, Srivastava appears to teaches away from employing past approaches of

load balancing (Jindal). As Srivastava teaches away from using past load balancing approaches, one skilled in the art would not look to combine Srivastava and Jindal. Thus the combination proposed by the Examiner is improper.

Accordingly, withdrawal of this rejection is respectfully requested. Claims 2-8, and 10-12 depend from claim 1 and are allowable for at least the same reasons as the independent claim from which they depend.

Independent Claim 15

Claim 15, as amended, recites a client for use in a client-server system. The client arranged to request a service, the request specifying a service address common to all of a plurality of servers in a server farm, each of the plurality of servers arranged to provide the service to the client and the servers communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm, connect to a first one of the plurality of servers in the server farm selected according to a load balancing method, receive the identity and status information relating to each of the servers in the server farm, from the first server in the server farm to which the client is connected, said information identifying each of the plurality of servers, and select a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information.

The Applicant respectfully submits that Jindal does not disclose the elements of claim 15, as amended. In particular Jindal does not teach or suggest a client arranged to “arranged to request a service, the request specifying a service address common to all of a plurality of servers in a server farm, each of the plurality of servers arranged to provide the service to the client and the servers communicating with one another so as to update identity and status information,” as recited in claim 15 as amended. As discussed above with respect to independent claim 1, Jindal discusses that “requiring the DNS server to query or access another server in order to resolve the request is inefficient and delays satisfaction of the request.” (Col. 2, lines 11-13). Thus, Jindal teaches away from the idea of storing information anywhere other than at the DNS server if at all possible, and does not teach or suggest that the information could be stored at a sever in the server farm. Jindal further teaches using only one data file for storing information collected on

the available servers. (Please see e.g. col. 8, lines 15-29). Therefore, Jindal could not reasonably teach or suggest storing a copy of the information at *each* of the servers in the servers farm.

In addition, Jindal teaches that once the client request is directed to the server and the client is connected to the sever that will provide the application, the process described by Jindal ends as a load balanced connection has been established to the application the client has requested. (Please see Abstract). Therefore, Jindal does not teach or suggest a client arranged to “receive the identity and status information relating to each of the servers in the server farm, from the first server in the server farm to which the client is connected, said information identifying each of the plurality of servers,” as recited in claim 15, as amended. Nor does Jindal teach or suggest a client arranged to “select a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information,” as recited in claim 15, as amended.

Srivastava fails to supply the missing limitations of claim 15, as amended. Srivastava does not describe a method or apparatus that allows the client to make any selection. Rather, Srivastava teaches that after selections are made at load balancing nodes, label records provide a mapping to fast switch packets. (please see Abstract and para. 0030). Although the background of Srivastava suggests that a client may perform a lookup at the DNS server, the background does not teach or suggest a client arranged to “request a service, the request specifying a service address common to all of a plurality of servers in a server farm, each of the plurality of servers arranged to provide the service to the client and the servers communicating with one another so as to update identity and status information,” as recited in claim 15 as amended. Nor does Srivastava teach or suggest a client arranged to “receive the identity and status information relating to each of the servers in the server farm, from the first server in the server farm to which the client is connected, said information identifying each of the plurality of servers,” as recited in claim 15, as amended. Lastly, Srivstava does not teach or suggest a client arranged to “select a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information,” as recited in claim 15, as amended. Therefore, the combination of Jindal and Srivastava (assuming, without admitting, that the combination of references was proper and feasible) would not render obvious that which is recited in claim 15.

As discussed above with respect to independent claim 1, Srivastava appears to teach away from the combination proposed. Srivastava criticizes conventional approaches of load balancing in the context of high-demand content network and criticizes past approaches for failing to provide “stickiness.” (See para. 0012 and 0024; please see also para 0004). Thus one skilled in the art would not be motivated to make the combination proposed in the Office Action, and, therefore, the proposed combination is improper.

The Applicant respectfully requests that the rejection be withdrawn. Claims 17-20 depend from claim 15 and are allowable for at least the same reasons as the independent claim from which they depend.

Independent Claim 21

Independent claim 21, as amended, recites a server for use in a client-server system having a plurality of servers in a server farm, each of the servers in the server farm being arranged to provide a service to the client, each of the servers being associated with a service address common to all of the servers, and the servers communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm. The server is arranged to connect to the client in response to a request from the client for the service routed to the server based on a load balancing method, the request specifying the common service address, send the identity and status information stored at the server to the client, and connect to the client in response to a selection, at the client, of one of the plurality of servers as the server to be used to provide the service to the client.

The Applicant respectfully submits that Jindal does not disclose the elements of claim 21 as amended. In particular, Jindal does not teach or suggest a server “having a plurality of servers in a server farm, each of the servers in the server farm being arranged to provide a service to the client, each of the servers being associated with a service address common to all of the servers, and the servers communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm,” as recited in claim 21, as amended. As discussed above with respect to independent claim 1, Jindal teaches that “requiring the DNS server to query or access another server in order to resolve the request is inefficient and delays satisfaction of the request.” (Col. 2, lines 11-13). Therefore, Jindal teaches away from the idea of storing information anywhere other than at the DNS server if at all

possible, and does not teach or suggest that the information could be stored at a sever in the server farm. Jindal further teaches using only one data file for storing information collected on the available servers. (Please see e.g. col. 8, lines 15-29). Therefore, Jindal could not reasonably teach or suggest storing a copy of the information at *each* of the servers in the servers farm.

In addition, Jindal teaches that once the client request is directed to the server and the client is connected to the sever that will provide the application, the process described by Jindal ends as a load balanced connection has been established to the application the client has requested. (Please see Abstract). Therefore, Jindal does not teach or suggest a server arranged to “send the identity and status information stored at the server to the client,” as recited in claim 21, as amended. Nor does Jindal teach or suggest a server arranged to “connect to the client in response to a selection, at the client, of one of the plurality of servers as the server to be used to provide the service to the client,” as recited in claim 21, as amended.

Srivastava fails to supply the missing limitations of claim 21, as amended. Srivastava does not describe a method or apparatus that allows the client to make any selection. Rather, Srivastava teaches that after selections are made at load balancing nodes, label records provide a mapping to fast switch packets. (please see Abstract and para. 0030). Although the background of Srivastava suggests that a client may perform a lookup at the DNS server, neither teaches or suggests a server “having a plurality of servers in a server farm, each of the servers in the server farm being arranged to provide a service to the client, each of the servers being associated with a service address common to all of the servers, and the servers communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm,” as recited in claim 21, as amended.

Nor does Srivastava teach or suggest a sever arranged to “send the identity and status information stored at the server to the client,” as recited in claim 21, as amended. Lastly, Srivstava does not teach or suggest a server arranged to “connect to the client in response to a selection, at the client, of one of the plurality of servers as the server to be used to provide the service to the client,” as recited in claim 21, as amended. Therefore, the combination of Jindal and Srivastava (assuming, without admitting, that the combination of references was proper and feasible) would not render obvious that which is recited in claim 21.

As discussed above with respect to independent claim 1, Srivastava appears to teach away from the combination proposed. Srivastava criticizes conventional approaches of load

balancing in the context of high-demand content network and criticizes past approaches for failing to provide “stickiness.” (See para. 0012 and 0024; please see also para 0004). Thus one skilled in the art would not be motivated to make the combination proposed in the Office Action, and, therefore, the proposed combination is improper.

The Applicant respectfully requests that the rejection be withdrawn. Claim 22 depends from claim 21 and is allowable for at least the same reasons as the independent claim from which it depends.

Independent Claim 23

Independent Claim 23, as amended, recites a client-server system having a plurality of servers in a server farm, each of the servers being arranged to provide a service to the client and each of the servers being associated with a service address common to all of the servers. The system is arranged to communicate information between the servers so that each of the plurality of servers maintains identity and status information relating to all of the servers, receive, at a DNS system, a request for the service from the client, the request specifying the common service address, apply a load balancing method to select a first one of the plurality of servers in the server farm and to connect the client to the first one of the plurality of servers in response to the request, send server information to the client from the server to which the client is connected, said server information identifying each of the plurality of servers and indicating the status of each of the plurality of servers to the client, and select, at the client, a second one of the plurality of servers as the server to be used to provide the service to the client, based on the server information.

The Applicant respectfully submits that Jindal does not disclose the elements of claim 23, as amended. In particular, Jindal does not teach or suggest a system arranged to “communicate information between the servers so that each of the plurality of servers maintains identity and status information relating to all of the servers,” as recited in claim 23, as amended. As discussed above with respect to independent claim 1, Jindal teaches that “requiring the DNS server to query or access another server in order to resolve the request is inefficient and delays satisfaction of the request.” (Col. 2, lines 11-13). Therefore, Jindal teaches away from the idea of storing information anywhere other than at the DNS server if at all possible, and does not teach or suggest that the information could be stored at a sever in the server farm. Jindal further

teaches using only one data file for storing information collected on the available servers. (Please see e.g. col. 8, lines 15-29). Therefore, Jindal could not reasonably teach or suggest storing a copy of the information at *each* of the servers in the servers farm.

In addition, Jindal teaches that once the client request is directed to the server and the client is connected to the sever that will provide the application, the process described by Jindal ends as a load balanced connection has been established to the application the client has requested. (Please see Abstract). Therefore, Jindal does not teach or suggest a system arranged to “send server information to the client from the server to which the client is connected, said server information identifying each of the plurality of servers and indicating the status of each of the plurality of servers to the client,” as recited in claim 23, as amended. Nor does Jindal teach or suggest a system arranged to “select, at the client, a second one of the plurality of servers as the server to be used to provide the service to the client, based on the server information,” as recited in claim 23, as amended.

Srivastava fails to supply the missing limitations of claim 23, as amended. Srivastava does not describe a method or apparatus that allows the client to make any selection. Rather, Srivastava teaches that after selections are made at load balancing nodes, label records provide a mapping to fast switch packets. (please see Abstract and para. 0030). Although Srivastava suggests in the background that a client may perform a lookup at the DNS server, it does not teach or suggest a system arranged to “communicate information between the servers so that each of the plurality of servers maintains identity and status information relating to all of the servers,” as recited in claim 23, as amended.

Nor does Srivastava teach or suggest a system arranged to “send server information to the client from the server to which the client is connected, said server information identifying each of the plurality of servers and indicating the status of each of the plurality of servers to the client,” as recited in claim 23, as amended. Lastly, Srivstava does not teach or suggest a system arranged to “select, at the client, a second one of the plurality of servers as the server to be used to provide the service to the client, based on the server information,” as recited in claim 23, as amended. Therefore, the combination of Jindal and Srivastava (assuming, without admitting, that the combination of references was proper and feasible) would not render obvious that which is recited in claim 21.

As discussed above with respect to independent claim 1, Srivastava appears to teach away from the combination proposed. Srivastava criticizes conventional approaches of load balancing in the context of high-demand content network and criticizes past approaches for failing to provide “stickiness.” (See para. 0012 and 0024; please see also para. 0004). Thus one skilled in the art would not be motivated to make the combination proposed in the Office Action, and, therefore, the proposed combination is improper.

The Applicant respectfully requests that the rejection be withdrawn. Claim 25-27 depend from claim 23 and are allowable for at least the same reasons as the independent claim from which they depends.

Dependent Claims 13-14

The Office Action rejected claims 13-14 under 35 U.S.C. §103(a) as being unpatentable over “Jindal” in view of “Srivastava” and further in view of U.S. Patent Application No. 2003/0149653 issued to Neill Penney, (hereinafter “Penney”).

As discussed with respect to independent claim 1 from which claim 13 and 14 depend, first, the combination of Jindal and Srivastava is improper, and second the combination does not teach or suggest that which is recited in independent claim 1, as amended. Claims 13 and 14 depend from claim 1 and are allowable for at least the reasons discussed above with respect to independent claim 1. Further, Penny does not supply the missing limitations discussed above with respect to claim 1.

Penny is directed to a method and apparatus for conducting financial transactions that enables a user to negotiate with multiple customers simultaneously, and receive and respond to transaction solicitations and amend requests in real time. (Abstract). Applicant notes that one skilled in the art would not seek to combine the load-balancing systems of Jindal and Srivastava with financial transaction manager of Penny. Thus, one skilled in the art would not be motivated to combine the references as suggested by the Examiner. Therefore, the conclusion that a client would be granted a second chance to connect to the desired server resulting in greater client satisfaction cannot be reached and is improper hindsight in light of the Applicant’s disclosure.

As discussed with respect to independent claim 1 from which 13 and 14 depend, the combination is improper, and even if combined do not teach that which is recited in independent claim 1. If Jindal/Srivastava were combined, there is no motivation to further modify the

Jindal/Srivastava combination with the teachings of Penny. Srivastava specifically notes that conventional approaches have disadvantages when applied in the context of high-demand content network. (para. 0012). Srivastava also criticizes approaches that do not provide for “stickiness.” (para. 0024). Thus Jindal/Srivastava combination would teach away from any additional combination with the conventional approaches of Penny. Moreover, one skilled in the art would not seek to modify the load-balancing system of Jindal/Srivastava with the financial transaction manager of Penny.

Even if the references were combined as suggested by the Examiner, the combination does not render obvious that which is recited in independent claim 1. In particular, none of the references teach or suggest servers in a server farm “communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended. Penny employs “conventional technology” for receiving and responding to solicitations to conduct or amend financial transactions. (Page 3, para. [0026]). Thus, Penny employs conventional technology that does not teach or suggest servers in a server farm “communicating with one another so as to update identity and status information stored at each of the servers relating to each of the servers in the server farm” as recited in claim 1, as amended.

In addition, none of the references teach or suggest “receiving, at the client, the identity and status information relating to each of the plurality of servers in the server farm, from the first server in the server farm to which the client is connected” as recited in claim 1, as amended. Nor do they teach or suggest “selecting, at the client, a second one of the plurality of servers in the server farm as the server to be used to provide the service to the client, based on the received information,” as recited in claim 1, as amended.

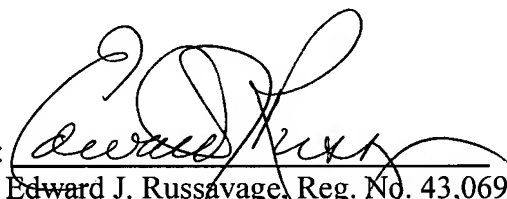
As the proposed combination is first improper, and even if combined (assuming without admitting proper and feasible) does not render obvious that which is recited in claim 1, as amended, the Applicant respectfully submits the rejection should be withdrawn. Claims 13 and 14 depend from claim 1 and are allowable for at least the same reasons. Accordingly, withdrawal of this rejection is respectfully requested.

CONCLUSION

In view of the foregoing amendments and remarks, reconsideration is respectfully requested. This application should now be in condition for allowance; a notice to this effect is respectfully requested. If the Examiner believes, after this amendment, that the application is not in condition for allowance, the Examiner is requested to call the Applicant's attorney at the telephone number listed below.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 50/2762.

Respectfully submitted,
Adam Stanley James Hawley, Applicant

By: 
Edward J. Russavage, Reg. No. 43,069
Matthew H. Grady, Reg. No. 52,957
LOWRIE, LANDO & ANASTASI, LLP
One Main Street
Cambridge, Massachusetts 02142
United States of America
Telephone: 617-395-7000
Facsimile: 617-395-7070

Docket No.: C2000-7000US

Date: June 14, 2007